

## STAR Computing at PDSF

E.L. Hjort, D.L. Olson, J. Yang

*Nuclear Science Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA*

The major resources for STAR computing are divided between two sites - the RHIC Computing Facility (RCF) at BNL, and PDSF, which is a part of NERSC at LBNL. In general the reconstruction of STAR data is done at RCF and PDSF is used primarily for production-level embedding and simulations. The bulk of the computing at PDSF is done in batch mode on its cluster of about 350 CPUs. Output files are typically archived in HPSS and then retrieved for further processing onto one of STAR's data vaults. The capacity of the data vaults is about 6TB for production data, 5TB assigned to STAR's physics working groups and 1TB for general scratch space. By the end of 2001 STAR had archived about 50TB in HPSS.

In 2001 STAR ran nine times as many jobs at PDSF as in 2000 and consumed five times the processing time. While this is in part a reflection of increases in PDSF computing capacity in 2001 (and STAR's need for this capacity) it is also a result of increased use of PDSF by institutions remote to LBNL. In 2000 59% of the jobs run at PDSF were submitted by LBNL users whereas in 2001 LBNL accounted for only 21% of the jobs run.

To support STAR computing at PDSF it is necessary to transfer files from RCF to PDSF. In particular the output from each reconstruction pass performed at RCF is replicated in its entirety at PDSF along with about 10% of the raw data files which are used to do embedding at PDSF. By using up to 20 parallel FTP sessions bandwidths approaching 10MB/s can be achieved. In practice transfers of up to 2TB/week are realized.

Embedding jobs are STAR's most resource-intensive use of PDSF. Embedding jobs are used to determine the efficiency of the STAR detector by generating simulated tracks and passing them through detector response simulators to generate simulated raw data. This data is then combined with real raw data and passed through the reconstruction chain. In 2001 STAR embedding jobs at PDSF reconstructed over 8 million events and generated about 11 TB of output. Efficiencies were calculated for particles including pi-, Omega, anti-Omega, Xi, anti-Xi, proton, anti-proton, Phi, K+, K-, Lambda, anti-Lambda, K0-short, K\*, anti-K\*, pi0 and photons.

A new resource for STAR computing at PDSF is called the high-bandwidth nodes. After testing with a prototype configuration of 10 nodes in 2000 the high bandwidth nodes have been expanded to 42 dual processor nodes, each node equipped with a local 300GB disk. These nodes are typically loaded with DST and uDST data and allow users to scan entire datasets with excellent I/O performance. At present the data on each disk are static and an interface ensures that each user's job is run on the node where the desired input file is locally resident.